

DEVICE AND PROCESS FOR FLOW CONTROL IN A
SWITCHED NETWORK

DISCLOSURE

5

Technical field

The present invention relates to a device and process for flow control in a switched network, notably in the field of communication between onboard equipment, such as for example computers, in the avionics field.

State of prior technology

Communication networks of known art used in the avionics field often have an architecture similar to that represented in figure 1, corresponding to norm ARINC 429, in which n devices C_1, C_2, \dots, C_n , for example computers, exchange information. Each transmitter has a physical link to each of the receivers. The information sent by one transmitter may occupy the entire bandwidth of the physical links concerned. Each of the devices C_1, C_2, \dots, C_n outputs a transmission line E_1, E_2, \dots, E_n linked to all the other devices which are considered as receivers in relation to this transmission line. This architecture has the advantage that it guarantees transmission of information from one transmitter to one or more receivers, without risk of conflict with the information sent by the other transmitters, since the latter use

different physical links. The period of communication from a transmitter to a receiver is thus defined deterministically. However, when the number of interconnected devices is increased, the number of lines of communication very rapidly becomes very large. An architecture of this kind then has many disadvantages in terms of weight, quantity of connections (hence potential risks of breakdown), complexity, wiring time, maintenance, etc.

There are also other methods of communication which are based on bus-type architectures, as illustrated in figure 2, corresponding notably to norms ARINC 629 and MIL 1553. Multiplexed information then circulates in "half-duplex" mode between the various devices C_1 , C_2 , ... C_n on a bus 10. This type of architecture is advantageous owing to the simplicity of the corresponding physical wiring. On the other hand, it requires that risks of collisions between the information from the various physical transmitters are managed, since the latter share the same physical link. In addition, the transmission rate is reduced compared to a point-to-point link since the bandwidth of each physical link is shared between different transmitters.

It is also known to use a star-shaped architecture around a central interface 20 implementing network switching functions, as illustrated in figure 3. n devices C_1 , C_2 , ... C_n are now linked to a switch 20 by

means of links L_1, L_2, \dots, L_n . The number of physical links (wire, etc.) is thus reduced.

More generally, the latter architecture may be
5 extended to a network containing several switches 20
linked in cascade fashion, as illustrated in figure 4.
It may be used in IT networks of the "Ethernet" type for
communication between terminals. A network of this kind
has the advantage that it requires only a relatively
10 small number of links when the number of interconnected
terminals grows large. However, a terminal C_i ($i=1, \dots, n$)
has only a single physical link, coming from the
corresponding switch 20, for receiving information from
all the other terminals and directed at this terminal C_i .
15 There is thus a risk of conflict if several terminals
attempt to send simultaneously a large quantity of
information to C_i . This problem is resolved by using the
notion of a virtual link: a virtual link is in effect a
logical link allowing information to be sent from a
20 transmitter to at least one receiver, with each virtual
link using at least one physical link.

As illustrated in figure 5, each virtual link is
single-direction from a transmitter to one or more
25 receivers. The information is sent in the form of
packages of data comprising a header representing the
virtual link number. In an "Ethernet" network switch 20
manages the virtual links and transmission rates
dynamically in order to adapt optimally to the

instantaneous network traffic. The bandwidths of the various virtual links are allocated such that at any time the sum of the bandwidths of all the virtual links using a given physical link is less than the theoretical
5 bandwidth of the said physical link, a bandwidth being a transmission capacity expressed as a number of items of information over a given unit of time, for example bit/s. Conversely, this dynamic allocation of bandwidths of the various virtual links does not enable it to be guaranteed
10 that an item of information will be transmitted between two terminals in a given time. This is a major disadvantage for applications requiring imposed frequencies of refreshment of exchanged data.

15 The purpose of the invention is to overcome these disadvantages by providing a device and process for controlling flows in a switched network.

Account of the invention

20

The invention relates to a device for controlling flows in a switched network comprising at least one transmitter device and at least one receiver device linked together across at least one switch, in which a
25 virtual link, which is a logical link using at least one physical link, enables information to be sent from a transmitter to at least one receiver, characterised in that each switch contains an allocation table (T), defined statically, which associates a bandwidth with each

of the virtual links so as to guarantee a maximum period for transmission of an item of information on each virtual link and an allocation such that for every physical link the sum of the bandwidths allocated to the various virtual links using this physical link is less than the bandwidth of this physical link.

Ideally the allocation table is such that a bandwidth may be allocated to a set of flows.

This device thus enables a maximum transmission period of an item of information on each virtual link to be guaranteed. To guarantee that there will never be any congestion of the communication network, this allocation is such that for every physical link, the sum of the bandwidths allocated to the various virtual links using this physical link is less than the bandwidth of this physical link, this bandwidth being dependent on the characteristics of the physical support.

The invention also relates to a flow control process in a switched network, comprising at least one transmitter device and at least one receiver device linked together across at least one switch, in which a virtual link, which is a logical link using at least one physical link, enables information to be sent from a transmitter to at least one receiver, characterised in that, in the switch, use is made of an allocation table (T), defined statically, which associates a bandwidth

with each of the virtual links so as to guarantee a maximum transmission time of an item of information in each virtual link and an allocation such that for every physical link the sum of the bandwidths allocated to the various virtual links using this physical link is less than the bandwidth of this physical link.

Ideally the allocation table is such that a bandwidth may be allocated to a set of flows.

Brief description of the drawings

Figure 1 illustrates a communication network of known art.

Figure 2 illustrates a bus-type architecture of known art.

Figures 3 to 5 illustrates star-shape architectures of known art.

Figure 6 illustrates the device of the invention.

Figure 7 illustrates an example of embodiment implementing the process of the invention.

Detailed account of embodiments

The device of the invention uses a physical layer of the "Ethernet" type, as illustrated in figure 6. It implements specific switches in which are defined, statically, an allocation table T the function of which

is to associate a bandwidth with each of the virtual links.

Such a table T has the following form:

| Virtual link | Physical ports | | Bandwidth |
|--------------|----------------|-----------|-----------|
| | Transmitters | Receivers | |
| . | | | |
| 14 | 1 | 3 | B_{14} |
| 15 | 1 | 2 | B_{15} |
| | | n | |
| 16 | 3 | 2 | B_{16} |
| . | | | |

5

This static allocation of the bandwidths to the various virtual links allows a maximum period of transmission of an item of information from a transmitter device to one or more receiver devices to be guaranteed.

10 Table T, defined above, thus has major advantages in terms of operational security.

In another embodiment, a bandwidth is allocated to a set of virtual links. When several virtual links, using
 15 at least one common physical link, are never all active simultaneously, it is possible to group these virtual links together. In the switch's table this group E of virtual links may then be allocated a smaller bandwidth than the sum of the bandwidths which would have been
 20 allocated to the various virtual links considered separately, without harming the performance of the device in respect of the maximum period of transmission of an item of information.

Since group E of virtual links requires a smaller bandwidth in the physical link, it is possible to pass more virtual links over this physical link. Or, if this is not necessary, it is possible to increase the bandwidth of this group of virtual links using the same physical link), enabling the maximum period of transmission of this information to be reduced.

Bearing in mind that the bandwidth required for a sub-set of virtual links may be less than the sum of the individual bandwidths of each of these links considered independently from one another. This embodiment consists in defining the switch's table such that it is possible to allocate a bandwidth to a group E of virtual links.

| Virtual link | Physical ports | | Bandwidth |
|---------------------------------|-----------------------|-----------------------|-----------|
| | Transmitters | Receivers | |
| . | | | |
| 0 | 0 | 1 2 3 . n | B_0 |
| 1 } 2 } 3 } E . n } | 1 2 3 . n | 0 | B_G |
| . | | | |

More generally, the virtual links sharing a given bandwidth may arrive at one or more switching ports.

The following case may be cited as an example: during aircraft stopovers in airports, it is sometimes necessary to download updates of equipment, or to upload information, from a station 21, which is itself linked to the network as illustrated in figure 7. An operation of this kind requires a virtual link VL_0 of the download station 21 to each of the devices C_1, C_2, \dots, C_n (transmission of information to the devices), and n virtual links VL_1, \dots, VL_n from the various devices to download station 21 (transmission of acknowledgement messages concerning the download). In practice the download is usually undertaken towards a single device at once. Consequently only one of the n virtual links $VL_1 \dots VL_n$ is used at a given time. However, since allocation table T is static, a bandwidth must be allocated to each of these n links. Since allocation table T is identical during the flight phases and the stopover phases, the bandwidth usable for downloading is relatively small (since it is above all necessary to guarantee good communication between the various devices during flight), although sufficient to ensure a satisfactory data transmission rate from download station 21 to the devices. However, in the opposite direction (from the devices to the download station), the n virtual links VL_1, \dots, VL_n must share a bandwidth of similar size to the previous one. The individual bandwidth of each of them is thus very small, the effect of which may be to slow down the download.

5
10
15
20